
Anisotropy profiling

IF TECHNIQUES such as neighbourhood tractography were to enable the robust segmentation of tracts representing equivalent fasciculi from a group of brain volumes, the question then arises, “What links or differentiates these tracts?” From a clinical perspective, we might be interested in looking for general differences in tract integrity between a healthy population and one affected by pathology. The work described in this chapter, which was completed under the supervision of Prof. David Laidlaw, attempts to look at integrity—as indicated by fractional anisotropy—on a fine-grained level, profiled along the length of a tract. The aim is to facilitate the testing of hypotheses about integrity at the within-tract level, and to investigate the behaviour and variability of anisotropy along a tract. This problem is separate to the one that neighbourhood tractography tries to solve, and is treated as such. We find evidence to suggest that although within-subject and within-group variance is large when FA is examined point-by-point, there can be sufficient regional differences between groups to ensure that subtle effects may well be masked by considering only mean FA values.

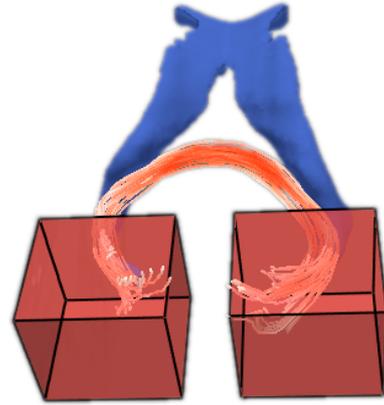
9.1 A single profile

To the extent that tractography is used at present for comparative clinical study, the most common approach is to average FA within the region segmented by the algorithm (e.g. Kanaan *et al.*, 2006), which may be represented by a line or a field. Region-averaged FA—however the region of interest is established—is a simple and useful way to study the effect of pathology on white matter integrity whilst controlling noise issues. On the other hand, ever greater numbers of studies are finding reduced FA effects in all kinds of pathologies, making such observations increasingly nonspecific; and since dMRI is the only available technique for studying structural white matter connectivity *in vivo*, independent corroboration or refutation of these results is extremely difficult. A partial list of scenarios in which reductions in FA have been observed could include schizophrenia (Ardekani *et al.*, 2003), multiple sclerosis (Ciccarelli *et al.*, 2003b), ischaemic leukoaraiosis in lacunar stroke (Jones *et al.*, 1999), epileptic patients after corpus callosotomy (Concha *et al.*, 2006), ischaemic stroke (Muñoz Maniega *et al.*, 2004) and normal ageing (O’Sullivan *et al.*, 2001).

For this study, six normal volunteers and five patients with vascular cognitive impairment (vci; a type of cognitive deficit which affects white matter and is quite common in elderly people) underwent a DTI protocol on a Siemens 1.5 T clinical scanner, with 12 noncollinear diffusion weighting gradient directions at a b -value of 1000 s mm^{-2} . The tractography infrastructure used for this work, BrainApp, uses a diffusion tensor-based deterministic streamlining algorithm, and visualises the results in terms of streamtubes and streamsurfaces (Zhang *et al.*, 2003). It uses whole brain seeding—which is possible in a reasonable time using a deterministic tractography algorithm—and thus avoids the selection constraint implicit to a neighbourhood or ROI seeding strategy. Simple streamline-based tractography lends itself very naturally to linear anisotropy profiling.

A streamtube is simply a piecewise linear streamline represented by a series of cylinders,

Figure 9.1: Example of a splenium streamtube set, segmented by placing a large region of interest near each end of the structure and retaining tubes passing through both. The shade of each streamtube indicates the local FA value. The blue structure represents the ventricles.



whose local radii may be constant or may be used to represent some characteristic of interest. A similar visualisation method has been used in other studies, such as Jones *et al.* (2005b). Working with tracts represented by single lines—rather than fields—is helpful for this work because it removes the need to linearise each tract before an anisotropy profile can be created. Ignoring its width, a streamtube, t_i , is therefore made up of piecewise linear line segments connecting a sequence of points, $(\mathbf{p}_{i,a})$, with $a \in \{1..N_i\}$, in the native acquisition space of the subject. The distance between successive points, d_i , is fixed in this space. Each of these tubes has a seed point, but unlike in the probability field output generated by `FSL ProbTrack`, the location of the seed point is not significant for the interpretation of the results, so we will not give it special treatment.

We first need to establish which tubes are of interest. Since BrainApp seeds throughout the brain, some kind of restriction is needed in order to focus on a specific white matter structure. Whatever method is chosen should be reproducible, however, so that it can be carried forward to comparative profiling between subjects. We used a two region of interest constraint to select the splenium of the corpus callosum, our tract of interest, with one ROI placed near the left end of the splenium tract and the other placed near the right end. These ROIs are symmetric, as per Conturo *et al.* (1999)—that is, they are treated identically, so swapping them would have no effect on the segmentation. This is not generally the case when one ROI provides the set of seed points, as in Abe *et al.* (2004) and some other studies. When working with streamtubes, this strategy amounts to taking the intersection of the set of tubes passing through ROI one, with the set of tubes passing through ROI two. An example of the result is shown in Fig. 9.1.

Since we have been critical of multiple ROI methods in earlier chapters, we will take a moment to justify this strategy. The important factors here are that the tractography algorithm being used to generate the streamtubes is deterministic, and that the seed points that generate the relevant streamtubes cannot be expected to form a compact neighbourhood, due to the whole brain seeding policy. Our objection about the effect of ROI constraints on interpretability due to the addition of extra conditional dependencies (cf. §6.5) only applies to output with a probabilistic significance. Constraining the algorithm by the selection of seed points is less relevant here; and neighbourhood tractography, which works on that principle, is not directly applicable. Ultimately, since the splenium is a coherent bundle with a distinctive shape, and is reasonably distinct from the rest of the corpus callosum and other nearby tracts in terms of the regions it connects together, the two ROI method is quite specific and reproducible enough. Moreover, it simply selects a set of streamtubes, just as choosing a number of seed points or clustering the streamtubes would. The effects are equivalent in essence.

Having “selected” the structure of interest, we can then plot the FA value, $f_{i,a}$, at each point on a tube, $\mathbf{p}_{i,a}$, thus forming an FA profile along the tube. (These values are interpolated from the FA data available at each voxel location.) Since all of these tubes are defined in the same space, aligning them is quite straightforward: we simply choose a plane which each streamtube must cross and consider the crossing points in each tube to be equivalent. We then examine the variability across the set of tubes at each point. This process produces a streamtube-averaged profile like the one shown in Fig. 9.2. In this case the distance between successive points, d_i ,

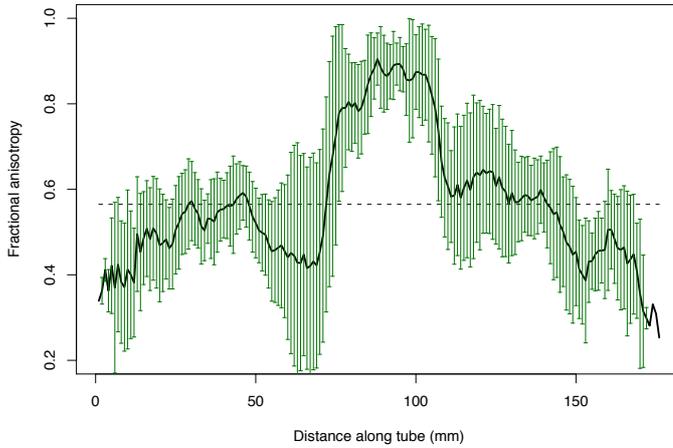


Figure 9.2: Pointwise mean FA along the set of splenium tubes segmented using the two ROI method in a single subject. The error bars indicate the mean plus or minus one standard deviation. The dashed line indicates the mean FA of the profile.

is 1 mm for all tubes. It should be noted that the alignment will handle differences in length well, but large shape differences, including kinks in some tubes, will render it inappropriate in some regions; and this is increasingly likely to occur as one moves away from the landmark plane.

Fig. 9.2 highlights two things in particular. Firstly, it is clear that the standard deviations are large to very large, relative to the means. Note, however, that on the left side of the graph in particular, the standard deviations are very large in a region near the middle of the tract, where the alignment plane was placed, and then shrink again further from the middle. This suggests that the variability is not primarily due to misalignment. One likely alternative cause is variation in the extent of partial volume effects. Some tubes will be nearer to the edge of the bundle than others, and the anisotropy at these locations is therefore more likely to be affected by proximity to grey matter or csf.

9.2 The median tube

Comparative profiling introduces some further issues. The questions of tube selection and alignment need to be reexamined, and differences in brain size must be compensated for in some way. We cannot simply use every tube selected in each brain, since the number of tubes selected is not fixed so bias would occur. We can't align tubes naïvely to a plane because each brain is represented in its own independent native space. And brain size cannot be neglected because it will affect the curvature of the structure and so the point homology.

Our approach to the first problem is to work only with the median tube from each brain; that is, the tube that minimises the average distance to all other tubes in the set. So, for a single subject whose splenium tube set contains N tubes, the median tube, t_m , is identified by calculating

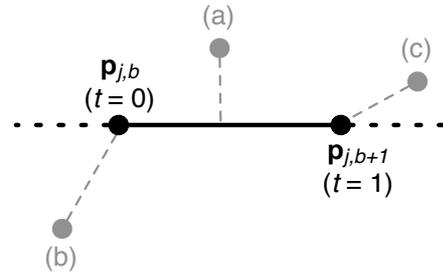
$$m = \arg \min_j \left\{ \frac{1}{N-1} \sum_{i=1; i \neq j}^N D(t_i, t_j) \right\}, \quad (9.1)$$

where $D(t_i, t_j)$ is the distance between streamtube i and streamtube j , given by the average distance from the points on the longer tube to the shorter tube, viz.

$$D(t_i, t_j) = \begin{cases} \frac{1}{N_i} \sum_{a=1}^{N_i} d(\mathbf{p}_{i,a}, t_j) & \text{if } N_i \geq N_j \\ \frac{1}{N_j} \sum_{b=1}^{N_j} d(\mathbf{p}_{j,b}, t_i) & \text{otherwise.} \end{cases}$$

The point-to-tube distance, $d(\mathbf{p}_{i,a}, t_j)$, is given by the minimum distance between the point and a line segment delimited by successive points in tube t_j . The point-to-segment distance, in

Figure 9.3: The calculation of the distance, \hat{d} , between a point and line segment depends on whether the projection of the point onto the segment direction crosses the segment itself. In (a), the projection crosses the segment ($0 < u < 1$), while in (b) and (c) it does not; and in these latter cases the shortest distances to the segment (dashed lines) are to one of its end points. The dotted extension of the line segment is shown for illustration.



turn, depends on the spatial arrangement of the point and segment (see Fig. 9.3). Mathematically, we parameterise the b th line segment as $\mathbf{s}_{j,b}(t) = \mathbf{p}_{j,b} + t\mathbf{l}_{j,b}$, where $\mathbf{l}_{j,b} = \mathbf{p}_{j,b+1} - \mathbf{p}_{j,b}$ and $t \in [0, 1]$. The projection of the point $\mathbf{p}_{i,a}$ onto the line segment—which forms the closest point between the two—is then given by $\mathbf{p}_{j,b} + u\mathbf{l}_{j,b}$, where

$$u = \frac{\mathbf{l}_{j,b} \cdot (\mathbf{p}_{i,a} - \mathbf{p}_{j,b})}{\mathbf{l}_{j,b} \cdot \mathbf{l}_{j,b}}. \quad (9.2)$$

The point and the line segment must, of course, be in the same space as one another. Now, $u \in \mathbb{R}$, and the distance between the point and the segment is calculated differently depending on whether the projection actually falls within the segment—i.e. $u \in (0, 1)$ —or not. Specifically,

$$\hat{d}(\mathbf{p}_{i,a}, \mathbf{s}_{j,b}) = \begin{cases} \|\mathbf{p}_{i,a} - \mathbf{p}_{j,b}\| & \text{if } u \leq 0 \\ \|\mathbf{p}_{i,a} - (\mathbf{p}_{j,b} + u\mathbf{l}_{j,b})\| & \text{if } 0 < u < 1 \\ \|\mathbf{p}_{i,a} - \mathbf{p}_{j,b+1}\| & \text{if } u \geq 1. \end{cases} \quad (9.3)$$

We then have

$$d(\mathbf{p}_{i,a}, t_j) = \min_b \{\hat{d}(\mathbf{p}_{i,a}, \mathbf{s}_{j,b})\}. \quad (9.4)$$

This is a standard formulation of the distance between a point and a line segment, but unfortunately it is a case in which the maths makes a simple concept look complicated. Eq. (9.2) is mathematical infrastructure for Eq. (9.3), which embodies the fact that if the line orthogonal to the line segment and passing through the point $\mathbf{p}_{i,a}$ does not cross the line segment, then the nearest point on the segment is in fact one of the end points. Fig. 9.3 illustrates this, for all three cases in Eq. (9.3). Note that if the next line segment, from $\mathbf{p}_{j,b+1}$ to $\mathbf{p}_{j,b+2}$, were to be collinear with the one illustrated, then the distance from point (c) to that segment would be lower than the distance shown, affecting the value of Eq. (9.4) appropriately.

Thus—finally—Eq. (9.1) is fully defined, and we can find the median tube in this way for each subject. This arrangement has the advantage that the median will tend to be towards the physical centre of a bundle of tubes, and therefore any partial volume effects should be relatively small. Incidentally, this justification differs slightly from that given for using the median streamline for tract matching in chapter 8, where the median was used simply because it epitomises the shape of a set of streamlines.

9.3 Intersubject tube alignment

As we have already mentioned, the tube sets representing the splenium of each subject's corpus callosum are necessarily each defined in their own space; and so absolute point locations are not directly comparable between subjects. In order to work around this complication, we observe that the splenium, being an interhemispheric fasciculus, always crosses the brain's midsagittal divide. (In fact, the placement of the rois guarantees this, since one is in the left hemisphere and one in the right.) This divide can be acceptably approximated by a plane. A number of methods have been proposed for automatically extracting this plane (e.g. Hu & Nowinski, 2003; Volkau *et al.*, 2006), but for this work we established its location in each subject manually, by placing four points, \mathbf{r}_1 to \mathbf{r}_4 , on the midsagittal divide by eye—thus marking the corners of a trapezium. Since three points are sufficient to establish a plane, the distance of the

fourth point to the plane was used as a simple error measurement to gauge the consistency of the placement. This distance is given by

$$\delta = \mathbf{n} \cdot (\mathbf{r}_4 - \mathbf{r}_1),$$

where

$$\mathbf{n} = \frac{(\mathbf{r}_2 - \mathbf{r}_1) \wedge (\mathbf{r}_3 - \mathbf{r}_1)}{\|(\mathbf{r}_2 - \mathbf{r}_1) \wedge (\mathbf{r}_3 - \mathbf{r}_1)\|},$$

the unit vector normal to the plane on which the points \mathbf{r}_1 to \mathbf{r}_3 lie. The mean placement error across all subjects, $\langle \delta \rangle$, was 0.90 mm.

Having established this midplane, we can find the location where each subject's median tube crosses the plane by first working out in which line segment the crossing occurs, and then finding the exact intersection of that segment with the plane. If the relevant line segment passes through the points \mathbf{r}_5 and \mathbf{r}_6 , it can be expressed as

$$\mathbf{s}(t) = \mathbf{r}_5 + t(\mathbf{r}_6 - \mathbf{r}_5),$$

and a bit more geometry yields the value of t where the line segment crosses the plane to be given by the ratio of determinants

$$t = - \frac{\det \begin{bmatrix} 1 & 1 & 1 & 1 \\ x_1 & x_2 & x_3 & x_5 \\ y_1 & y_2 & y_3 & y_5 \\ z_1 & z_2 & z_3 & z_5 \end{bmatrix}}{\det \begin{bmatrix} 1 & 1 & 1 & 0 \\ x_1 & x_2 & x_3 & x_6 - x_5 \\ y_1 & y_2 & y_3 & y_6 - y_5 \\ z_1 & z_2 & z_3 & z_6 - z_5 \end{bmatrix}},$$

where $\mathbf{r}_1 = (x_1, y_1, z_1)$ and so on. We then translate the co-ordinate system of each native space so that this intersection point is at the origin. Finally, on the assumption that the point where the median tube crosses the midsagittal divide is equivalent across brains, we treat all of these translated spaces as being equivalent. It is now possible to combine the median tubes from all subjects into an *intersubject* tube set, and find an intersubject median tube from this set.

Correcting for translational differences between subjects is not sufficient, however, since the shapes of the different subjects' spleniums will still vary due to differences in brain size. One approach to this problem is to use the intersubject median-of-medians tube, t_M , as a spatial reference, and take an f_A value, f' , for each tube at each point on this median by finding the nearest neighbour point on each separate tube. That is,

$$f'_{i,a} = f_{i,\tilde{b}},$$

where

$$\tilde{b} = \arg \min_b \|\mathbf{p}_{M,a} - \mathbf{p}_{i,b}\|;$$

so the a th f_A value from tube i is the value at that point on t_i that is closest to the a th point on t_M , with i now indexing over subjects. This gives us a one-dimensional f_A profile of fixed length for all subjects.

9.4 Comparative profiling

Fig. 9.4 shows the result of performing the whole process described above on a full data set. We located the splenium, using the two ROI method, in each subject. We then calculated a single intersubject median tube by combining all subjects' individual medians together; but subsequently separated them into patient and control groups once more for generating the averaged profiles shown in the figure. The intersubject median's f_A data was not included,

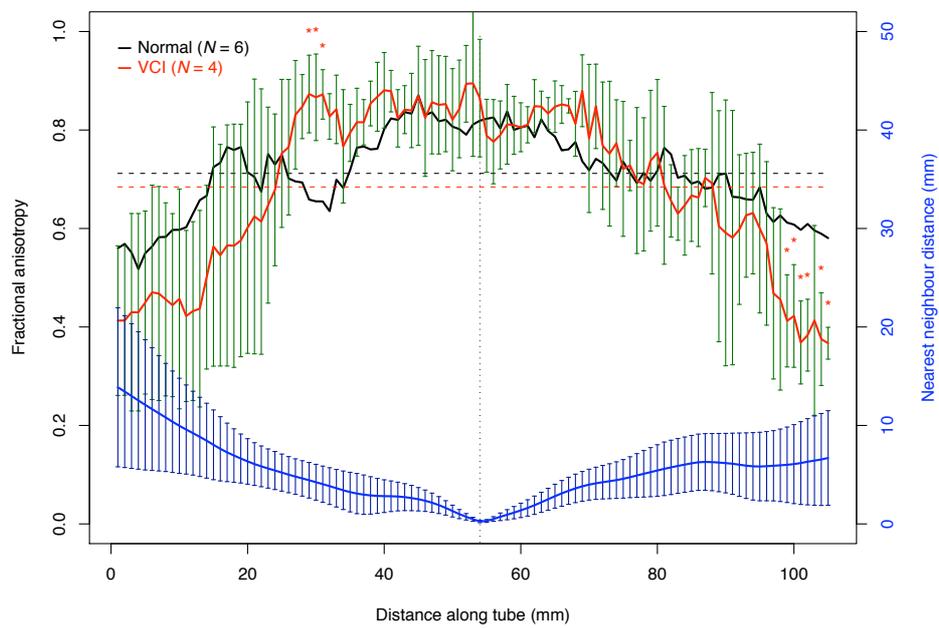


Figure 9.4: An example of comparative profiling between groups of subjects. The red line with green error bars shows the average (plus or minus one standard deviation) value of f' , averaged across all subjects with VCI, at each point on the intersubject median tube. The black line shows the mean across the normal subjects. Appropriately coloured horizontal dashed lines show the profile mean FA. The blue line with blue error bars indicates the mean (plus or minus one standard deviation) distance from the intersubject median tube to its nearest neighbour at each point, across all subjects. The vertical dotted line shows the location of the midplane.

reducing the number of vCI subjects contributing to four. These initial results were first presented in Clayden *et al.* (2007b).

In Fig. 9.4, red stars indicate significant ($P < 0.05$) differences between the groups using a two tailed t -test on f' data at each point. Since no correction for multiple comparisons was performed, these differences are tentative results at best, but they are somewhat informative nevertheless. Since the significant points are clustered into two (almost) contiguous regions, it seems unlikely that the differences are due to random noise effects; although the combination of interpolation and the nearest neighbour process makes successive points somewhat interdependent. It is interesting to note that while the grand mean FA , indicated by dashed horizontal lines, is lower for the vCI group than for the normal group—although this difference was not significant—the two regions differing between the profiles are not consistent in the sign of the difference between the groups. The region at the right hand end of Fig. 9.4 shows lower FA in the vCI group, which is the most common finding in pathological groups, while the region on the left side of the graph shows *higher* FA in vCI. This may be because the region represents an area of crossing fibres. If one of the two fibre populations were to preferentially suffer a loss of integrity, an increase in FA would be expected. To the left of the significant region, FA is decreased relative to the normal population again, although the error bars are too large for this to be significant.

The large nearest neighbour distances in this latter region may be responsible for the large variability which is particularly noticeable at the left hand end of the profiles. The blue curve indicates the mean and standard deviation of the distances from the intersubject median tube to each subject's individual median. This is zero by definition at the midplane—indicated with a vertical dotted line—and tends to increase as one gets further from there. The greater this distance becomes, the greater the divergence of the median tubes from one another; but it is not clear whether, or to what extent, an increase in divergence makes the profiled FA values intrinsically less comparable.

9.5 Discussion

The approach to anisotropy profiling described above has allowed us to explore some of the issues involved with this kind of comparative analysis, and to get a sense of the variability in anisotropy along a major tract. There are, however, evident reasons that this technique would not be very widely useful in its current state. Firstly, not all tracts in the brain cross the midsagittal divide, so using this landmark for intersubject alignment will not be possible in all cases. Secondly, the use of nearest neighbours for establishing a point homology between tubes is not robust, and the performance of the technique will depend on the shape of the tract of interest. One possible way of avoiding both of these issues is to use registration for alignment of median tubes between subjects. This would solve the problem of handling differences in brain size at the same time as annulling translational misalignment. It would be less tract-specific than the combination of techniques described in §9.3, and so if it worked well enough it would be applicable, in theory, to any tract of interest in the brain. Another possible avenue would be to use the probabilistic neighbourhood tractography methods described in chapter 8 to select a representative line for each subject, rather than taking the median. This would circumvent the limitations of the two ROI method in more complex tracts than the splenium.

In addition to dealing with these systematic limitations, we would need to apply the profiling process to more data to get a clearer picture of its effectiveness, or to draw any serious clinical conclusions. In particular, it would be interesting to study differences in the profiles between scans of a single subject, and between two normal populations. We would also need to look at other tracts. It may be that the full FA profile is actually too noisy a representation to be generally useful; but it is nevertheless suggested by the results so far that the mean FA along a streamtube, or group of tubes, is only a perfunctory summary of the available information. Fig. 9.2 shows that even allowing for large error bounds the FA along a tract in a single subject is not well encapsulated by the mean, and Fig. 9.4 demonstrates potential regions of difference between healthy and possibly abnormal profiles despite there being no significant difference in the means. We have also done some work in which the profiles were filtered for high

frequency noise by convolving them with a Gaussian smoothing kernel—which has not been shown here—but it remains unclear whether or not this would be beneficial. It may be that one could use this kind of smoothing to make multiple comparisons correction less of a problem, as *vBM* does, but the choice of variance for the smoothing kernel might be hard to justify. All of this is left as future work.

The ability to meaningfully compare anisotropy—or diffusivity, or any other measure of interest—between groups at a fine-grained but tract-specific level could be very useful for comparative analysis in white matter, but for the moment there are, as we have discussed, a number of hurdles in the way of the profiling approach we have described here.